arXiv:0905.2386v3 [cs.DM] 6 Aug 2010

Combinatorial information distance

Joel Ratsaby

Department of Electrical and Electronics Engineering, Ariel University Center
Ariel 40700, ISRAEL
ratsaby@ariel.ac.il

Summary. *Let $\#A$ denote the cardinality of a finite set A. Let $t(x) = x$ if $x \geq 1$ and 1 otherwise. For any two sets A, B denote by $\delta(A,B) = \log_2\left(t\left(\#\left(B \cap \overline{A}\right)\#A\right)\right)$. We define a new set distance $d(A,B) = \max\left\{\delta\left(A,B\right), \delta\left(B,A\right)\right\}$ motivated by combinatorial notions of entropy and information [8]. We prove that d is a semi-metric on the space of sets of size at least 2. The triangle inequality. holds for triplets A, B, C that are not strictly contained one in another.*

Key words: *Distance function, Lempel-Ziv complexity, Binary sequences*

1 Introduction

A basic problem in pattern recognition [6] is to find a numerical value that represents the dissimilarity or 'distance' between any two input patterns of the domain. For instance, between two binary sequences that represent document files or between genetic sequences of two living organisms. There are many distances defined in different fields of mathematics, engineering and computer and information sciences [5]. A good distance is one which picks out only the 'true' dissimilarities and ignores those that arise from irrelevant attributes or due to noise. In most applications the design of a good distance requires inside information about the domain, for instance, in the field of information retrieval [4] the distance between two documents is weighted largely by words that appear less frequently since the words which appear more frequently are less informative. The ubiquitous Levenshtein-distance [9] measures the distance between two sequences (strings) as the minimal number of edits (insertion, deletion or substitution of a single character) needed to transform one string into another. Approximate string matching [10] is an area that uses such edit-distances to find matches for short strings inside long texts. Typically, different domains require the design of different distance functions which take such specific prior knowledge into account. It can therefore be an expensive process to acquire expertise in order to formulate a good distance. The paper of [20] introduced a notion of complexity of finite binary string which does not require any prior knowledge about the domain or context represented by the string (this is sometimes referred to as the *universal* property). This complexity (called the production complexity of a

string) is defined as the minimal number of copy-operations needed to produce the string from a starting short-string called the base. This definition of complexity is related to Levenshtein-distance mentioned above. It is proportional to the number of distinct phrases and the rate of their occurrence along the sequence. There has been some work on using the LZ-complexity to define a sequence-distance measure in bioinformatics [17]. Other applications of the LZ-complexity include: approximate matching of strings [16], analysis of complexity of biomedical signals [2], recognition of structural regularities [11], characterization of DNA sequences [7] and responses of neurons to different stimuli [3], study of brain function [18] and brain information transmission [19] and EEG complexity in patients [1].

In the current paper we introduce a distance function between two strings which also possesses this universal property. Our approach is to consider a binary string as a *set* of substrings [14]. To represent the complexity of such a set we use the notion of combinatorial entropy [12] and introduce a new set distance function. We proceed to describe some fundamental concepts concerning entropy and information of sets.

2 Entropy and information of a set

Kolmogorov [8] investigated a non-stochastic measure of information for an object y. Here y is taken to be any element in a finite space $\mathbb{Y}$ of objects. He defines the 'entropy' of $\mathbb{Y}$ as $H(\mathbb{Y}) = \log \#\mathbb{Y}$ where $\#\mathbb{Y}$ denotes the cardinality of $\mathbb{Y}$ and all logarithms henceforth are taken with respect to 2.

As he writes, if it is known that $\mathbb{Y} = \{y\}$ then this provides $\log \#\mathbb{Y}$ bits of 'information' or in his words "this much entropy is eliminated". To represent partial information about $\mathbb{Y}$ based on another information source $\mathbb{X}$ let $R = \mathbb{X} \times \mathbb{Y}$ be a general finite domain and consider a set

$$A \subseteq R \tag{1}$$

that consists of all permissible pairs $(x, y) \in R$ (in the usual probabilistic-based representation of information this is analogous to having a uniform prior probability distribution over a certain region of the domain). The entropy of $\mathbb{Y}$ is defined as

$$H(\mathbb{Y}) = \log \#\Pi_{\mathbb{Y}}(A)$$

where $\Pi_{\mathbb{Y}}(A) \equiv \{y \in \mathbb{Y} : (x, y) \in A \text{ for some } x \in \mathbb{X}\}$ denotes the projection of A on $\mathbb{Y}$. Consider the restriction of A on $\mathbb{Y}$ based on x which is defined as

$$Y_x = \{y \in \mathbb{Y} : (x, y) \in A\},\ x \in \Pi_{\mathbb{X}}(A) \tag{2}$$

then the conditional combinatorial entropy of $\mathbb{Y}$ given x is defined as

$$H(\mathbb{Y}|x) = \log \#Y_x. \tag{3}$$

Kolmogorov defines the information conveyed by x about $\mathbb{Y}$ by the quantity

$$I(x : \mathbb{Y}) = H(\mathbb{Y}) - H(\mathbb{Y}|x). \tag{4}$$

In [15] an alternative view of $I(x : \mathbb{Y})$ is defined as the information that a set Y_x conveys about another set $\mathbb{Y}$ satisfying $Y_x \subseteq \mathbb{Y}$. Here the domain R is defined based

on the previous set A as $R = \Pi_{\mathbb{Y}}(A) \times \Pi_{\mathbb{Y}}(A)$ which consists of all permissible pairs (y, y') of objects. Knowledge of $x \in \mathbb{X}$ means knowing the set $A_x \subseteq R$, $A_x = \{(y, y') : y \in \Pi_{\mathbb{Y}}(A), y' \in Y_x\}$. The information between Y_x and $\mathbb{Y}$ is then defined as

$$\begin{aligned} I(Y_x : \mathbb{Y}) &= \log \left(\#\Pi_{\mathbb{Y}}(A)\right)^2 - \log \#A_x \\ &= \log \left(\#\Pi_{\mathbb{Y}}(A)\right)^2 - \log(\#\Pi_{\mathbb{Y}}(A)\#Y_x). \end{aligned} \tag{5}$$

Clearly, $I(Y_x : \mathbb{Y}) = I(x : \mathbb{Y})$. Note that $I(Y_x : \mathbb{Y})$ measures the difference in description length of any *pair* of objects $(y, y') \in \Pi_{\mathbb{Y}}(A) \times \Pi_{\mathbb{Y}}(A)$ when no 'labeling' information exists versus that when there exists information which labels one of them as being an element of Y_x. Thus the second term in (5) can be viewed as the conditional combinatorial entropy of $\Pi_{\mathbb{Y}}(A)$ given the set Y_x. In [12, 15, 13] this is used to extend Kolmogorov's combinatorial information to a more general setting where knowledge of x still leaves some vagueness about the possible value of y.

While the distance that we introduce in this paper is general enough for any objects, our interest is to introduce a combinatorial distance for binary strings. We henceforth refer to $\mathbb{X} = \{0, 1\}^*$ as the space of binary strings x. Each string $x \in \mathbb{X}$ is a *description* of a corresponding set Y_x in the space $\mathbb{Y}$ of objects y. Our approach to defining a distance between two binary strings x and x' is to relate them to sets of objects and then measure the distance between the two corresponding sets. Let us denote by $M : \mathbb{X} \to \mathbb{Y}$ the function which defines how a string x yields a set $Y_x \subseteq \mathbb{Y}$. In general, M may be a many-to-one function since there may be several strings (viewed as descriptions of the set) of different lengths for a given set. In the context of the above, we now consider a permissible pair $(x, y) \in A$ to be one which consists of an object y that is contained in a set Y_x which is described by x. Clearly, not every possible pair (x, y) is permissible, as for instance, if $y' \notin Y_x$ then (x, y') is not permissible.

In the next section we introduce a combinatorial information distance. We start with a distance for general sets and then apply it as a distance between binary strings.

3 The distance

Let Ω be a domain. For a finite set $A \subset \Omega$ denote by $\#A$ the cardinality of A. The cardinality of the empty set is zero. Define the following function:

$$t(x) = \begin{cases} x & \text{if } x \geq 1 \\ 1 & \text{otherwise .} \end{cases}$$

Definition 1. For any two finite sets $A, B \subset \Omega$ define the following function $\delta : \Omega \times \Omega \to \mathbb{N}_0$ which maps a pair of finite sets into the non-negative integers:

$$\delta(A, B) = \log \left(t \left(\# \left(B \cap \overline{A}\right) \#A\right)\right)$$

where $\overline{A}$ denotes the complement of the set A. It is simple to realize that $\delta(A, B)$ equals $\log \left(\# \left(B \cap \overline{A}\right) \#A\right)$ with the exception when A or B are empty or $B \subseteq A$.

Remark 2. Note that δ is non-symmetric, i.e., $\delta(A, B)$ is not necessarily equal to $\delta(B, A)$. Also, $\delta(A, B) = 0$ when $B \subseteq A$ (not only when $A = B$).

The definition of δ resembles in functional form the second log term in (5) and may therefore be interpreted as conditional combinatorial entropy of B given A. From an information theoretical perspective ([20]) the value $\log \#\left(B \cap \overline{A}\right)$ represents the additional description length (in bits) of an element in B given *a priori* knowledge of the set A. Hence we may view A as a partial 'dictionary' while the part of B that is not included in A takes an additional $\log \#\left(B \cap \overline{A}\right)$ bits of description given A.

The following set will serve as the underlying space on which we will consider our distance function. It is defined as

$$2^{\Omega}_{+} = 2^{\Omega} \setminus \{A \subset \Omega : \#A \leq 1\}.$$

It is the power set of Ω but without the empty set and singletons. We note that in practice for most domains, as for instance the domain of binary strings considered later, the restriction to sets of size greater than 1 is minor.

We have the following auxiliary lemma which will be useful in the proof of Theorem 5.

Lemma 3. *The function δ satisfies the triangle inequality on any three elements A, B, $C \in 2^{\Omega}_{+}$ none of which is strictly contained in the other.*

Proof. Suppose A, B, C are any elements of 2^{Ω}_{+} satisfying the given condition. Without loss of generality we will show that

$$\delta(A, C) \leq \delta(A, B) + \delta(B, C). \tag{6}$$

First we consider the specific case where the triplet has an identical pair. If $A = C$ then by Remark 2 it follows that $\delta(A, C) = 0$ which is a trivial lower bound so (6) holds. If $A = B$ then $\delta(A, B) = 0$ and both sides of (6) are equal hence the inequality holds (similarly for the case of $B = C$). Next we consider the more general case where each of the following three quantities is at least 1,

$$\#\left(C \cap \overline{A}\right),\ \#\left(B \cap \overline{A}\right),\ \#\left(C \cap \overline{B}\right) \geq 1. \tag{7}$$

By definition of 2^{Ω}_{+} we have $\#A \geq 2$ hence

$$\delta(A, C) = \log\left(t(\#(C \cap \overline{A})\#A)\right) = \log\left(\#(C \cap \overline{A})\#A\right) = \log \#\left(C \cap \overline{A}\right) + \log \#A.$$

Next, we claim that $C \cap \overline{A} \subseteq \left(B \cap \overline{A}\right) \cup \left(C \cap \overline{B}\right)$. Suppose $x \in C \cap \overline{A}$ then $x \in C$ and $x \in \overline{A}$. Now, either $x \in B$ or $x \in \overline{B}$. Suppose $x \in B$ then because $x \in \overline{A}$ it follows that $x \in B \cap \overline{A}$. Suppose $x \in \overline{B}$ then because $x \in C$ it follows that $x \in C \cap \overline{B}$. This proves the claim. Next, we have

$$\delta(A, B) + \delta(B, C) = \log \#A + \log \#(B \cap \overline{A}) + \log \#B + \log \#(C \cap \overline{B}).$$

It suffices to show that

$$\log \#\left(C \cap \overline{A}\right) \leq \log \#\left(B \cap \overline{A}\right) + \log \#\left(C \cap \overline{B}\right) + \log \#B. \tag{8}$$

We claim that if three non-empty sets X, Y, Z satisfy $X \subseteq Y \cup Z$ then $\log \#X \leq \log(2\#Y\#Z)$. To prove this, it suffices to show that $\#X \leq 2\#Y\#Z$. From the

given, we have $\#X \leq \#(Y \cup Z) \leq \#Y + \#Z$ and the following inequality holds for non-empty sets Y and Z:

$$\frac{1}{2 - {}^{1}/_{\#Y}} \leq \frac{1}{2-1} = 1 \leq \#Z.$$

Hence $\#Y \leq \#Z\,(2\#Y - 1)$. Therefore combining the above we have

$$\#X \leq \#Y + \#Z \leq 2\#Y\#Z$$

from which the claim follows. By (7), we may let $X = C \cap \overline{A}$, $Y = B \cap \overline{A}$ and $Z = C \cap \overline{B}$ and from both of the claims it follows that

$$\#\left(C \cap \overline{A}\right) \leq 2\#\left(B \cap \overline{A}\right)\#\left(C \cap \overline{B}\right). \tag{9}$$

Taking the log on both sides of (9) and using the inequality $2 \leq \#B$ (which follows from $B \in 2^{\Omega}_{+}$) we obtain

$$\log\#(C\cap\overline{A}) \leq 1+\log\#(B\cap\overline{A})+\log\#(C\cap\overline{B}) \leq \log\#B+\log\#(B\cap\overline{A})+\log\#(C\cap\overline{B}).$$

This proves (8). □ □

Next, we define the combinatorial information-distance.

Definition 4. For any two non-empty sets A, B define the *combinatorial information-distance* as

$$d(A,B) = \max\{\delta(A,B), \delta(B,A)\}.$$

In the following result we show that d satisfies the properties of a semi-metric.

Theorem 5. *The distance function d is a semi-metric on 2^{Ω}_{+}. It satisfies the triangle inequality for any triplet $A, B, C \in 2^{\Omega}_{+}$ such that no element in the triplet is strictly contained in another.*

Proof. Clearly d is symmetric as $d(A,B) = d(B,A)$. It is also clear that it is non-negative. From Remark 2 it is clear that for $A = B$, $\delta(A,B) = \delta(B,A) = 0$ hence $d(A,B) = 0$. For A, $B \in 2^{\Omega}_{+}$ such that $A \neq B$ (possibly $B \subset A$ or $A \subset B$) then $\delta(A,B) > 0$ or $\delta(B,A) > 0$ hence $d(A,B) > 0$. Hence d is a semi-metric on 2^{Ω}_{+}.

Next, we show that it satisfies the triangle inequality for any triplet $A, B, C \in 2^{\Omega}_{+}$ such that no element is strictly contained in another. For any non-negative numbers a_1, a_2, a_3, b_1, b_2, b_3, that satisfy

$$\begin{aligned} a_1 &\leq a_2 + a_3 \\ b_1 &\leq b_1 + b_2, \end{aligned} \tag{10}$$

we have

$$\begin{aligned} \max\{a_1, b_1\} &\leq \max\{a_2 + a_3, b_2 + b_3\} \\ &\leq \max\{\max\{a_2, b_2\} + \max\{a_3, b_3\}, \max\{b_2, a_2\} + \max\{b_3, a_3\}\} \\ &= \max\{a_2, b_2\} + \max\{a_3, b_3\}. \end{aligned}$$

From Lemma 2 it follows that (10) holds for the following: $a_1 = \delta(A,C)$, $b_1 = \delta(C,A)$, $a_2 = \delta(A,B)$, $b_2 = \delta(B,A)$, $a_3 = \delta(B,C)$, $b_3 = \delta(C,B)$. This yields

$$d(A,C) \leq d(A,B) + d(B,C)$$

hence d satisfies the triangle inequality for such a triplet. □ □

Let us now define the distance between two binary strings x and x'. We take as $\mathbb{X}$ the space of all finite binary strings and use the concepts and definitions introduced in section 2.

Definition 6. Let $R = \mathbb{X} \times \mathbb{Y}$ be all possible pairs (x,y) and let $A \subseteq R$ be the set of permissible pairs. For any $x \in \Pi_{\mathbb{X}}(A)$ denote by $Y_x = \{y \in \mathbb{Y} : (x,y) \in A\}$. Let $x, x' \in \mathbb{X}$ be two binary strings. Then the combinatorial information distance between x and x' is defined as

$$d(x,x') = d(Y_x, Y_{x'})$$

where $d(Y_x, Y_{x'})$ is defined in Definition 4.

The next result follows directly from Theorem 5.

Corollary 7. *Let $\mathbb{Y}$ be a space of objects y and $\mathbb{X}$ the space of binary strings x describing all sets $Y_x \subseteq \mathbb{Y}$ that have cardinality at least 2. The combinatorial information distance $d(x,x')$ is a semi-metric on $\mathbb{X}$ and satisfies the triangle inequality for triplets x, x',x'' whose sets Y_x, $Y_{x'}$, $Y_{x''}$ are not strictly contained one in another.*

As an example, consider the mapping M that takes binary strings to sets Y in $\mathbb{Y} = \{0,1\}^k$ (the k-cube) for some fixed finite k. This resembles the method of [20] who break up a binary string s into a set of substrings whose cardinality is taken to be the complexity of s. Consider the following scheme for describing a set Y: we form a string from the concatenation of all vertices on the k-cube that are elements of the set Y. For instance, suppose $k = 5$ and $Y = \{01100, 10101, 11110\}$ then any of the 6 possible strings that are formed by concatenating the three elements (we call them k-words) of the set in any order represent a possible description of Y. If a string has N repeating k-words then clearly only a single copy of these k-words will be placed in Y_x. Note that the mapping M that takes x to Y_x eliminates redundancy in a way that is similar to the method of [20] which gives the minimal number of copy operations needed to reproduce a string from a set of its substrings.

Another possible mapping M may be defined by scanning a fixed window of length k across the string x and collecting each substring (captured in the window) as an element of the generated set Y_x. It requires some empirical analysis on real data sets to determine the optimal value of the parameter k that yields a good distance. Yet another approach which does not need to choose k is to use the method of [20] and collect substrings of x (of possibly different lengths) as the set Y_x.

Whichever is the mapping, to compute the combinatorial information distance between any two finite strings x and x' first determine the sets of substrings for x and x' and let them be Y_x and $Y_{x'}$ respectively. The distance $d(x,x')$ according to Definition 6 is the set distance $d(Y_x, Y_{x'})$. Its properties are described in Corollary 7.

References

1. D. Abasolo, R. Hornero, C. Gomez, M. Garcia, and M. Lopez. Analysis of EEG background activity in Alzheimer's disease patients with Lempel-Ziv complexity and central tendency measure. *Med. Eng. Phys.*, 28(4):315–322, 2006.
2. M. Aboy, R. Hornero, D. Abasolo, and D. Alvarez. Interpretation of the Lempel-Ziv complexity measure in the context of biomedical signal analysis. *IEEE Trans. on Biomedical Eng.*, 53(11):2282–2287, 2006.
3. J. M. Amigo, J. Szczepaski, E. Wajnryb, and M. V. Sanchez-Vives. Estimating the entropy rate of spike trains via Lempel-Ziv complexity. *Neural Computation*, 16(4):717–736, 2004.
4. R. Baeza-Yates and B. Ribeiro-Neto. *Modern Information Retrieval*. Addison-Wesley, 1999.
5. M. Deza and E. Deza. *Encyclopedia of Distances*, volume 15 of *Series in Computer Science*. Springer-Verlag, 2009.
6. R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification*. Wiley-Interscience Publication, 2000.
7. V. D. Gusev and L. A. Nemytikova. On the complexity measures of genetic sequences. *Bioinformatics*, 15(12):994–999, 1999.
8. A. N. Kolmogorov. Three approaches to the quantitative definition of information. *Problems of Information Transmission*, 1:1–17, 1965.
9. V. Levenshtein. Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics Doklady*, 10:707–710, 1966.
10. G. Navarro. A guided tour to approximate string matching. *ACM Computing Surveys*, 33(1):31–88, 2001.
11. Y. L. Orlov and V. N. Potapov. Complexity: an Internet resource for analysis of DNA sequence complexity. *Nucleic Acids Research*, 32:W628–W633, 2004.
12. J. Ratsaby. On the combinatorial representation of information. In Danny Z. Chen and D. T. Lee, editors, *The Twelfth Annual International Computing and Combinatorics Conference (COCOON'06)*, volume LNCS 4112, pages 479–488. Springer-Verlag, 2006.
13. J. Ratsaby. Information efficiency. In *The* 33^{rd} *International Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM'07)*, volume LNCS 4362, pages 475–487. Springer-Verlag, 2007.
14. J. Ratsaby. A distance measure for properties of boolean functions. Presented at *Workshop on Boolean Functions: Theory, Algorithms and Application*, CRI, Haifa January 27 - February 1, 2008.
15. J. Ratsaby. Information width. *Technical Report # arXiv:0801.4790v2*, 2008.
16. L. Russo, G. Navarro, and A. Oliveira. Parallel and distributed compressed indexes. In *Proc. 21th Annual Symposium on Combinatorial Pattern Matching (CPM)*, LNCS 6129, pages 348–360, 2010.
17. K. Sayood and H. H. Otu. A new sequence distance measure for phylogenetic tree construction. *Bioinformatics*, 19(16):2122–2130, 2003.
18. X. Wu and J. Xu. Complexity and brain function. *Acta Biophysica Sinica*, 7:103–106, 1991.
19. J. Xu, Z. Liu, and R. Liu. Information transformation in human cerebral cortex. *Physica D*, 106:363–374, 1997.
20. J. Ziv and A. Lempel. On the complexity of finite sequences. *IEEE Transactions on Information Theory*, 22(3):75–81, 1976.